

# **CA ARCserve Replication and High Availability**

## **Debugging Guide r16.5**



This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the "Documentation") is for your informational purposes only and is subject to change or withdrawal by CA at any time. This Documentation may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This Documentation is confidential and proprietary information of CA and may not be disclosed by you or used for any purpose other than as may be permitted in (i) a separate agreement between you and CA governing your use of the CA software to which the Documentation relates; or (ii) a separate confidentiality agreement between you and CA.

Notwithstanding the foregoing, if you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all CA copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to CA that all copies and partial copies of the Documentation have been returned to CA or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS DOCUMENTATION "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL CA BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF CA IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice.

The manufacturer of this Documentation is CA.

Provided with "Restricted Rights." Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

Copyright © 2013 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

# Contents

<b>Chapter 1: Log Files .....</b>	<b>4</b>
<b>Chapter 2: Configuration and Trace Files .....</b>	<b>6</b>
<b>Chapter 3: Keywords to Use in Log Analysis .....</b>	<b>11</b>
<b>Chapter 4: Debugging Installation Errors .....</b>	<b>14</b>
<b>Chapter 5: Debugging Server Connection Errors .....</b>	<b>16</b>
<b>Chapter 6: Frequent Connection Loss and Reset .....</b>	<b>19</b>
<b>Chapter 7: Slow Replication Performance .....</b>	<b>20</b>
<b>Chapter 8: Debugging Driver Errors .....</b>	<b>24</b>

# Chapter1: Log Files

- Control Service Logs

- Default Location(s):

- C:\Program Files\CA\ARCserve RHA\Manager\log (on x86 Windows)
    - C:\Program Files (x86)\CA\ARCserve RHA\Manager\log (on x64 Windows)

- Names / Description

- **Ws\_man.log** – This is the current log that is being written to by the Control Service
    - **Ws\_man.log.#** - These are the logs that have been rotated
    - By default, the last 20 logs are saved.
    - Logs are rotated when a log file reaches 5 MB in size.

- RHA Engine Logs

- Default Location(s):

- C:\Program Files\CA\ARCserve RHA\Engine\log

- Names / Description

- **Ws\_rep.log** – This is the current log that is being written to by the Engine Service
    - **Ws\_rep.log.#** - These are the logs that have been rotated.
    - By default, the last 20 logs are saved. The number of logs to save is controlled by TraceNumberOfFilesToKeep parameter in ws\_rep.trc file located in the Engine installation directory (by default, C:\Program Files\CA\ARCserve RHA\Engine).
    - Logs are rotated when a log file reaches 5 MB in size.
    - **Syst\_stats.log** – This log file contains system information and it is written when the engine was started. This content is also added to the ws\_rep.log files when the engine is started.
    - **Ws\_rep.cfg** – This is a copy of the configuration file from when the engine was started.
    - **Run\_info\_<scenarioID>\_[1 | 2].log** – This is a copy of the scenario (1-from the Master, 2-from the Replica). The time stamp of the file shows when the scenario was last started.
    - **Run\_info\_<scenarioID+1>\_[1 | 2].log** – This is a copy of the backward HA scenario (1-from the active server, 2-from the standby server). The time stamp of the file shows when the backward scenario was last started.

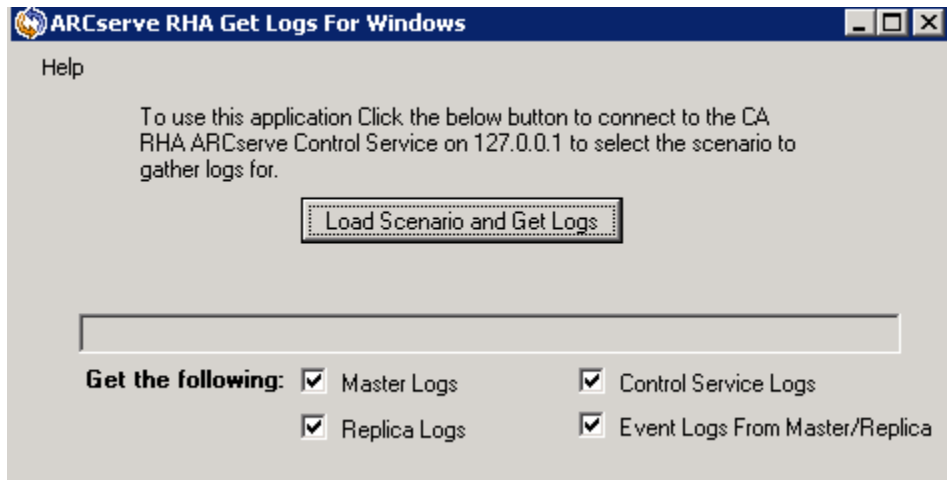
- **AFUtil\_<date>\_<time>#.log** – This is a specialized log that is used by the AFUtility.exe application in CA ARCserve D2D integrated scenarios.
  - **Ws\_p2v.log** – This is a specialized log used by CA Full System HA P2V service.
  - **Ws\_helper.log** and **ws\_helper.log.#** - These logs are used by CA ARCserve RHA Helper service in Exchange 2007 and Exchange 2010 scenarios. They contain Exchange Powershell commands invoked by RHA Engine and the associated error messages.
- **Powershell Snap-in Logs**
  - Default Location(s):
    - C:\Program Files\CA\ARCserve RHA\Powershell Snapin\log (on x86 Windows)
    - C:\Program Files (x86)\CA\ARCserve RHA\Powershell Snapin\log (on x64 Windows)
  - Names / Description
    - **xopowershell.log** – The most current log file that is being written to by the Powershell snap-in.
    - **xopowershell.log.#** - These are the rotated logs.
- **RHA Engine Installation logs**
  - Default Location: C:\CA\_Install\_log
  - Install-log.txt just tells when the logging was started and stopped
  - Install-utl2.log information from when the engine remote installation (captures only errors and memory issues by default)
  - To get a manual verbose output of the installer run the Engine installation file from a command line:
    - For 64-bit Windows:  
**CAARCserveRHAEngine64.exe "/V /log "c:\temp\caxoinstall.log"**
    - For 32-bit Windows:  
**CAARCserveRHAEngine.exe "/V /log "c:\temp\caxoinstall.log"**
    - The installation log is created in c:\temp\caxoinstall.log. Review the log file and search for errors.
    - On Windows 2008 use a command line with the elevated privilege. Open cmd as Administrator.
- **Collecting the logs**

To collect the Control Service logs and RHA Engine logs from all servers of a scenario use ws\_getlog.exe utility located on the Control Service host in [INSTALLDIR]\ws\_getlog folder. When running ws\_getlog.exe you have an option to get

- Master logs

- Replica logs
- Control Service logs
- Windows Application and System event logs from Master/Replica

By default, all the options are checked



Click on "Load Scenario and Get Logs" button". You'll be prompted to enter user account to authenticate with the Control Service. Once authenticated it pops up a window with a list of all scenarios. Select a scenario and click OK button. The application connects to all servers and downloads the requested logs zipping them in <scenario-name>.zip file.

**Note: To connect to remote servers ws\_getlog.exe uses RPC ports. If RPC port is not open on a remote server or the connection is blocked by a firewall the application won't be able to download the logs from the server. In this case you'll have to connect to the server, zip/compress RHA log folder, Windows application/system event logs and copy the zip files manually.**

## Chapter2: Configuration and Trace Files

- RHA Engine
  - Default installation folder
    - C:\Program Files\CA\ARCserve RHA\Engine
  - Names / Description
    - ws\_rep.exe – RHA Engine .
    - ws\_rep.cfg – RHA Engine configuration file.
    - ws\_rep.trc – RHA Engine trace file.
  - Ws\_rep.cfg Configuration Parameters

- The default value and the description of RHA Engine parameters is shown.
- # sign at the beginning of the line means that the line is commented out.
- After making changes to the file CA ARCServe RHA Engine service must be restarted to apply the changes
- In most cases you will not need to change any of the default settings. Change parameter values only if instructed by CA Support.

– Ws\_rep.trc Tracing Parameters

- The default tracing options are

```
TraceLogFileSize 5242880
TraceNumberOfFilesToKeep 20
PrintFunction off
PrintTime on
DefaultDebugLevel + debug
```

- TraceLogFileSize - the maximum size of the log file before the log is rotated.
- TraceNumberOfFilesToKeep – the number of log files to keep.
- LogDirectoryName – full path to the log folder.
- To enable trace to all function calls set  
PrintFunction=on  
DoIndentation=true.  
**Note: Enabling PrintFunction option increases the amount of data written to RHA log and makes logs harder to read. Enable it only when requested by CA Support**
- DefaultDebugLevel controls the trace level. The default is set to "DefaultDebugLevel + debug" writing ERROR, WARNING and DEBUG messages to RHA log.
- Additional trace options can be added to DefaultDebugLevel:  
DefaultDebugLevel + debug + <trace-level>
- Here is a list of useful tracing options when debugging specific errors:

+ **note** – this trace level is useful when debugging replication errors with certain files. It prints to the log all I/O operates in replicated files showing the name of the updated file, the size and the offset the written data chunk. In the following example the name of the file is /data/test1.txt, the size of I/O is 91 byte and the offset is 0:

```
(3/29/2013 7:55:24 PM) note-> Event Open, header 43, data 0, id 7, txf id 0, name /data1/test1.txt
(3/29/2013 7:55:24 PM) note-> Event Write, header 63, data 91, id 8, txf id 0, name /data1/test1.txt, offs
0, size 91, tm 1364601324
(3/29/2013 7:55:24 PM) note->Event Close, header 51, data 0, id 9, txf id 0, name /data1/test1.txt, last tm
1364601324
```

+note trace level can be enabled on the Master or Replica or on both servers.

- + **user1** – this trace level is use used for debugging connection problems with a protected application. This trace level is used mostly when
  - debugging problems with auto-discovery (must be enabled on the Master)
  - debugging problems in IsAlive check (must be enabled on the active server)
- + **user2** – this trace level is used mostly for debugging problems with XONET network protocol used for communication between all RHA Engines and the Control Service. When +user2 trace level is enabled an information about every sent XONET packet and every received acknowledgement (ACK) is written to RHA log. It includes
  - XONET packet ID
  - XONET packet length
  - File ID or command ID or ACK ID
  - destination IP

Master log:

```
== Thread 684 (CommandSender to calabsrv06-fvm1: 2384352659) ==
(4/1/2013 4:52:30 PM) user2->
Send Comm: h_id=1611148775, pack_len=148, app_id=8; rm.IP=10.130.65.126
== Thread 932 (XONET Thread Pool_1) ==
(4/1/2013 4:52:30 PM) user2->
Received CACK: h_id=1611148775, pack_len=3299, app_id=8; rm.IP=10.130.65.126
```

Replica log:

```
== Thread 2636 (XONET Duplex: 0 from 10.130.65.125) ==
(4/1/2013 4:52:30 PM) user2->
Received Comm: h_id=1611148775, pack_len=148, app_id=8; rm.IP=10.130.65.125
(4/1/2013 4:52:30 PM) user2->
Send CACK in pool: h_id=1611148775, pack_len=3299, app_id=8; rm.IP=10.130.65.125
```

- + **user3** - this trace level is used mostly when
  - debugging problems with script execution. Enabling this option returns error code from every script executed by RHA.
  - debugging problems with Block Synchronization. Enabling this trace level provides more logging information from a Block Synchronization.

**Note: Changing the trace level may significantly increase the amount of data written to RHA log and may impact replication performance. Logs rotate much faster and recent log data is quickly overwritten. Enable additional trace options only when debugging errors or when requested by CA Support.**

After ws\_rep.trc file is updated and saved there is no need to restart CA ARCServe RHA Engine service. The RHA Engine detects that the file was changed and reparses the file. It rotates the log starting a new ws\_rep.log. All trace parameters are printed to the log. The new log file starts with a message:

<DATE> TRC Default Trace config file exe\_dir/exe\_name.trc is to be reparsed:

The trace settings are written to the log. For example, if you enable +note you should see at the beginning of the log

*DefaultLoggingRank fatal +error +warning +debug +note*

If the RHA Engine log has not rotated or you don't see in ws\_rep.log the configured trace level, restart CA ARCServe RHA Engine service to apply changes in ws\_rep.trc.

## • Control Service

- Default installation folder

- C:\Program Files\CA\ARCServe RHA\Manager (on x86 Windows)
  - C:\Program Files (x86)\CA\ARCServe RHA\Manager (on x64 Windows)
- Names / Description
- ws\_man.exe – the Control Service executable .
  - mng\_core\_com.dll - the Control Service management DLL.
  - mng\_core\_com.cfg – the Control Service configuration file.
  - ws\_rep.trc – the Control Service trace file.
- mng\_core\_com.cfg Configuration Parameters
- Shows the default value and the description of the Control Service parameters.
  - # sign at the beginning of the line means that the line is commented out.
  - After making changes to the file CA ARCServe RHA Control Service must be restarted to apply the changes
  - Most of the parameters in mng\_core\_com.cfg do not apply to the Control Service. The specific parameters that apply to the Control Service are listed under "application WS\_MAN parameters"
  - In most cases you won't need to change any of the default settings. Change parameter values only if instructed by CA Support.
- ws\_man.trc Tracing Parameters
- The default tracing options are
 

```
TraceLogFileSize 5242880
TraceNumberOfFilesToKeep 20
PrintFunction off
PrintTime on
DefaultDebugLevel + debug
```
  - TraceLogFileSize - the maximum size of the log file before the log is rotated.
  - TraceNumberOfFilesToKeep – the number of log files to keep.
  - LogDirectoryName – full path to the log folder.
  - DefaultDebugLevel controls the trace level. The default is set to "DefaultDebugLevel + debug" writing ERROR, WARNING and DEBUG messages to RHA log.
  - Additional trace options can be added to DefaultDebugLevel:
 

```
DefaultDebugLevel + debug + <trace-level>
```
  - Here is a list of useful tracing options when debugging specific errors:

+ **note** – this tracing level is useful when debugging timeout issue with the Control Service command “Get scenario state” (usually when the Control Service manages 10-s of 100-s of scenarios)

+ **user1** – this trace level is useful for debugging the Control Service license enforcement issues and errors

+ **user2** – this trace level is used mostly for debugging problems with XONET network protocol in communication between the Control Service and RHA Engines. For details see RHA Engine tracing options.

+ **user3** - this trace level is used mostly when debugging thread deadlocks in the Control Service management DLL mng\_core\_com.dll

**Note: Changing the trace level increases the amount of data written to the Control Service log. Logs rotate much faster and recent log data is quickly overwritten. Enable additional trace options only when debugging errors or when requested by CA Support.**

- NAT Utility

- Default installation folder
  - C:\Program Files\CA\ARCserve RHA\Engine
- Names / Description
  - natutlgui.exe – GUI-based NAT configuration utility
  - natutl.exe – command line NAT configuration utility
- Configuration File
  - [INSTALLDIR]\nat\_setting.xml

For more information about RHA NAT Utility, see “Configure the NAT Utility” in CA ARCserve RHA Administration Guide.

## Chapter 3: Keywords to Use in Log Analysis

- **Error events** - have trace prefix "**error->**"

Example:

*(3/18/2013 5:15:01 PM) error->BlockSynch: Unable to remove path E:/Exchsrvr/mdbdata/Exchange Resources Storage Group/Public Folder Store/Catalog/pub000B13C0/Build/Indexer/CiFiles/00010003.ci, err=The process cannot access the file because it is being used by another process.(32)*

When debugging error messages it is important to check the error code returned by OS. In this example, the error code is 32 that means:

*The process cannot access the file because it is being used by another process*

This error is from a Replica log and it means that the file is in use by another process, therefore RHA Engine cannot remove it. When you see such an error you need to identify the process using the file and stop it. If you cannot find the process you can use a Microsoft Handle utility to identify the process using the file.

<http://technet.microsoft.com/en-us/sysinternals/bb896655.aspx>

A complete list of Windows system error codes can be found at

[http://msdn.microsoft.com/en-us/library/windows/desktop/ms681381\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/ms681381(v=vs.85).aspx)

**Note: In most cases when failing to execute an I/O operation that may impact integrity of data on a Replica RHA writes an error message to the log and retries to execute the operation again. If all retries fail then a notification error event is sent to the parent node in the replication tree, it is written to the Master log, sent to the Control Service and displayed in GUI event window.**

Read below about RHA notification events.

- **Notification events** – start with "**Notifying scenario** <scenario-ID>"

All scenario notification events start with **Notifying scenario** <scenario-ID>. These events are sent to the Control Service and displayed in GUI event window. They are also stored in the Control Service [INSTALLDIR]\ws\_events folder.

All notification events generated by any server in a scenario are always written to the Master log and sent to the Control Service and GUI. However, by default, RHA keeps only 20 log files. Old notification events may not be found in the Master logs if the associated log file has been overwritten. The Control Service stores notification events from every scenario in [INSTALLDIR]\ws\_events\<scenario-ID>\_scenario.txt file. When the Control Service is restarted it reloads these files and scenario events are not lost.

**Note: If notification events are not shown in RHA Engine logs because the logs have been overwritten you can get a list of notification events from the Control Service server [INSTALLDIR]\ws\_events\<scenarioID>\_scenario.txt file. By default, these events are also written to the application event log on the Master server**

If a notification event is sent by a Replica it is first written to the Replica log and then sent to the parent node in the replication tree and forwarded to the Control Service and GUI.

There are 5 main types of notification events:

- I - Informational
- S - Significant
- W - Warning
- E - Error
- C - Critical Error

You can see the type of the event from the first letter of the event ID. For example, in the notification event below:

*(4/2/2013 6:33:36 PM) Notifying scenario 2384351819, host 2; message id: ER00048, message: Unable to execute event Rename for C:/data1/d1.tmp : The process cannot access the file because it is being used by another process.(32)*

ER00048 is event ID. Prefix E means that this is an Error event.

**Note: A complete list of RHA Notification events is included in mess.txt file located in the Control Service and the Engine installation folders.**

When interpreting notification events always check the host ID this event was sent from. For example, in the above example, the event was sent by host 2 that means that this event was sent by a Replica. A Master server has always host ID = 1. Replica servers have host IDs starting from 2 incrementing by 1. If a scenario has more than one Replica server you can see host ID of a server from the name of **run\_info\_<scenarioID>\_<index>.log** in the log folder. Index shows the host ID of the server in the scenario. You can also see host ID of a particular server if you open the scenario file or run\_info log and search for "ReplNode"

`<object label="ReplNode" id="5000">`

Under it you'll see the host name, management and replication IP addresses, port number and the host "Index" that is the host ID

The Master server has host ID = 1:

`<data label="Index" val="1" type="Number"/>`

Continue searching for "ReplNode" and you'll find all hosts configured in the scenario file. For Replica servers Index value start with 2,

`<data label="Index" val="2" type="Number"/>`

- **Warning events** - have trace prefix **"warning->"**

These are non-critical warning events from the RHA Engine or OS.

- **Trace level events**

When additional trace level is enabled (+note, +user1, +user2, +user3) RHA starts logging all events linked to this trace level adding prefix *->trace-level*, for example:

->note  
->user1  
->user2  
->user3

- **FilterSendMessage events** – start with "FilterSendMessage <Operation> failed"

FilterSendMessage is used for sending messages from the XOMF filter driver. For example,

*80070070.*

*(4/1/2013 7:09:09 AM) FilterSendMessage XOMF\_RT\_JRN\_OP failed, result:  
There is not enough space on the disk.*

When FilterSendMessage event is a critical event RHA sends an associated notification event, for example,

*(4/1/2013 7:10:39 AM) Notifying scenario 232333228, host 1; message id:  
CD00379, message: Write to journal file failed (c000007f).*

and the scenario stops.

- **Invoke psscript** (applies to Exchange scenarios only).

Invoke psscript messages in ws\_helper log show the actual powershell (PS) command executed by ws\_helper service on an Exchange server. For example, the following command checks state of Exchange databases

*Information -> Invoke psscript \$sg = Get-StorageGroup -server CALABSRV05-EX2 | Where  
{(((\$\_.Name -eq 'Third Storage Group') -or (\$\_.DistinguishedName -eq 'Third Storage Group')) -and  
(\$\_.Recovery -ne \$True)}  
\$getmdbcmd = "Get-MailboxDatabase -StorageGroup '\$sg' -status"  
invoke-expression \$getmdbcmd | where { \$\_.Mounted -eq \$False }*

The executed PS commands here are

*Get-StorageGroup -server CALABSRV05-EX2 | Where {(((\$\_.Name -eq 'Third Storage Group') -or  
(\$\_.DistinguishedName -eq 'Third Storage Group')) -and (\$\_.Recovery -ne \$True)}  
\$getmdbcmd | where { \$\_.Mounted -eq False }*

These particular PS commands check if Exchange databases from 'Third Storage Group' are mounted on the server CALABSRV05-EX2. If the executed PS command returns an error there will be an associated error message in ws\_helper.log (it will have "Error->" prefix). You can try running this command manually in Exchange PS on this server and check why it fails.

## Chapter 4: Debugging Installation Errors

If installation of CA ARCserveRHA Engine fails with errors copy the installation file to the server where you are trying to install and run from a command line.

For 64-bit Windows:

**CAARCserveRHAEngine64.exe "/V /log "c:\temp\caxoinstall.log"**

For 32-bit Windows:

**CAARCserveRHAEngine.exe "/V /log "c:\temp\caxoinstall.log"**

The installation log is created in c:\temp\caxoinstall.log. Review the log file and search for errors.

One of the most common installation errors when installing or uninstalling CA ARCserveRHA Engine or upgrading it from one version to another is error 1720:

### **Error Message:**

**"Error 1720: There is a problem with this Windows Installer package. A script required for this install to complete could not be run. Contact your support personnel or package vendor. "**

Usually this error indicates that WMI repository is corrupted on the server.

Actions to try:

**Action1:** Re-register all of the dlls and in the wbem folder, re-registering WMI Service and Provider and restart WMI management service.

```
net stop winmgmt
```

```
cd %windir%\System32\Wbem
```

```
for %i in (*.dll) do RegSvr32 -s %i
```

```
net start winmgmt
```

**Action2:** if re-registering DLLs does not help and you are trying to install on a Windows 2008 or later run the following command:

**winmgmt /salvagerepository**

This command takes the content of the inconsistent repository and merges it into the rebuilt repository if it is readable (supported starting from Windows 2008). On Windows Server 2003 /salvagerepository switch is not supported and you must rebuild WMI repository manually (see Action4).

**Action3:** If the above doesn't work, then run:

**winmgmt /resetrepository**

This command resets WMI repository to its initial state when the OS was first installed. It is supported only starting from Windows 2008.

**Action4:** If neither of the above steps resolves the issue (or you are trying to install on a Windows 2003 where /salvagerepository or /resetrepository switches are not supported) rebuild WMI repository using the following steps:

1. Change startup type to Window Management Instrumentation (WMI) Service to disabled
2. Stop the WMI Service; you may need to stop IP Helper Service first or other dependent services before it allows you to stop WMI Service
3. Rename the repository folder: C:\WINDOWS\system32\wbem\Repository to Repository.old
4. Open a CMD Prompt with elevated privileges
5. cd %windir%\System32\Wbem
6. Re-register all DLLs from wbem folder  
for %i in (\*.dll) do RegSvr32 -s %i
7. Set the WMI Service type back to Automatic and start WMI Service
8. Recompile all localized descriptions of classes in .MOF and .MFL files in wbem folder  
for %i in (\*.mof, \*.mfl) do Mofcomp %i
9. Reboot the server.

**Note:** In certain cases installation fails due to lack rights in the service logon account. You can try installing CA ARCserveRHA Engine in the Local System account and change the account later.

# Chapter 5: Debugging Server Connection Errors

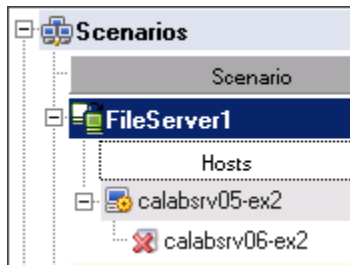
Scenario fails to start with error

*ER00069 Unable to connect to host <server-name>*

*EM00166 Unable to connect to <server-name>; the connection refused by engine*

*EM00166 Unable to connect to server; Attempt to connect, but timed out without establishing a connection*

*EM00166 Unable to connect to server; connection reset*



## Troubleshooting:

- Under scenario management check the status of the server. Red x indicates that the server is disconnected. If the server is shown as disconnected highlight the server in the **Scenarios** window and check the server properties (highlight the disconnected server click on properties tab). Check Host Connection properties and make sure that it has the correct IP address.

**Note: If a server is a Microsoft Cluster make sure that the IP address configured in the Host Properties is a virtual IP of the replicated source group and not a physical IP of the cluster node.**

- Check if the server is pingable (if ICMP ping requests are not blocked by Firewall)
- Remote into the server which had status disconnected in scenario. Go to Windows services and check if CA ARCserve RHA Engine service is running. If CA ARCserve RHA engine service is not running, then try starting it.
- Check that CA ARCserve RHA engine process ws\_rep.exe is listening on port 25000. Open a command line and run

```
netstat -ab > netstat.log
```

Open netstat.log in the Notepad and search for ws\_rep.exe. If ws\_rep.exe process is listening for incoming connections you should see the following line:

```
TCP    0.0.0.0:25000      <server-name>:0    LISTENING
[ws_rep.exe]
```

- If CA ARCserve RHA engine service is running and the process is listening on port 25000, check if the port 25000 is open in a Firewall and it is accessible from other servers. If Windows Firewall is enabled check in the firewall configuration that ws\_rep.exe process is allowed to communicate through the firewall and the port 25000 is open.

To check that the port is accessible from other servers open command prompt on another server and type the following command

```
telnet < IP address> 25000
```

use the IP address configured in the Host properties of the scenario.

**Note: On Windows 2008 Telnet client is not installed by default. To install it go to Server Manager, Features, Add Feature, select Telnet Client and proceed with installation instructions.**

Check that telnet can connect to this server port 25000. If telnet cannot connect when running it from a remote server, run this command locally on the server itself and check if telnet can connect to port 25000. If telnet connects to port 25000 when running locally on the server but it cannot connect when running telnet from a remote host the connection is either blocked by a firewall or there is a routing problem to this server. If you configured port forwarding on the firewall, check that it is properly configured and works.

- If a Master server is shown as disconnected and telnet running from the Control Service host can connect to port 25000 of the disconnected server, try restarting CA ARCserve RHA Control Service. After that close and reopen GUI, logon the Control Service and check the server status again. When restarted the Control Service reconnects to the Master servers from all scenarios and updates their state. When the Control Service connects to the RHA engine you'll should see in the engine log ws\_rep.log the associated handshaking message:

*XO handshake with <CS IP address> : 3889 Chunk=65536 Build #16.5.0.3409, passed successfully*

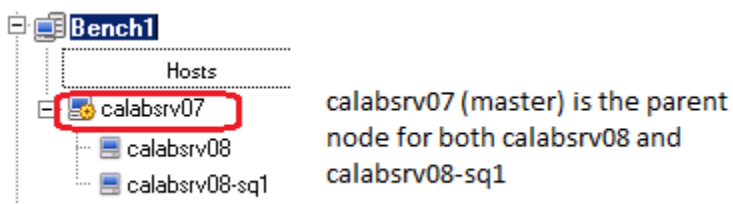
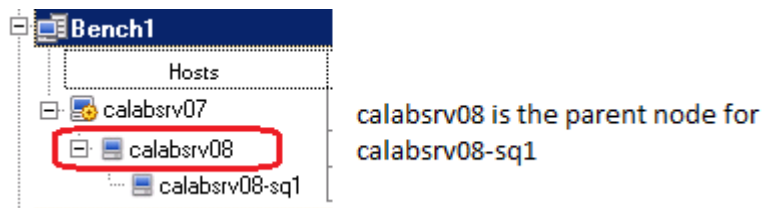
If you don't see the handshaking message in the engine log it means that the Control Service cannot connect to the RHA engine.

- Try restarting CA ARCserve RHA Engine service on the disconnected server. Prior to restarting the service remove config\_25000 folder located in the engine installation folder. The default path is C:\Program Files\CA\ARCserve RHA\Engine\config\_25000

**Note: on a Microsoft Cluster config\_25000 folder is located on shared drives of the replicated resource group <drive-letter>\ CA ARCserve RHA\config\_25000**

Stop CA ARCserve RHA service and check in Windows task manager that there are no zombie ws\_rep.exe processes running. Start CA ARCserve RHA service and check that it is listening on port 25000. Verify in GUI that the connection is restored.

- If the disconnected server is a Replica and you already tried restarting CA ARCserve RHA service on the server and you also confirmed that telnet client running from the server parent node can connect to port 25000, try restarting CA ARCserve RHA Engine service on the parent node. The parent node is a server one level above in the scenario replication tree. The following diagrams show 2 examples:



Remove config\_25000 folder on the parent node and restart CA ARCserve RHA service. Check that it is listening on port 25000 and verify in GUI that the connection is restored.

- Check the version of CA ARCserve RHA engine file. Go to CA ARCserve RHA installation folder, by default "C:\Program Files\CA\ARCserve RHA\Engine ", and check the version of ws\_rep.exe file. Verify that that the version of the file matches the version of the file on other servers. Check that it matches the version of the Control Service. To check the version of the Control Service in CA ARCserve RHA Manager go to Help → About.
- If neither of the above steps help to restore the connection try to reboot the server or call CA Support.

## Chapter 6: Frequent Connection Loss and Resets

When a running scenario sends a lot of connection loss and connection reset errors with a Replica server

*ER00067 Connection to host <server-name> lost*

*EX01030 Connection reset*

first validate that CA ARCserve RHA engine service is running on the Replica server and telnet client from the Master can connect Replica's port 25000. Open a command line on Master and run

```
telnet <Replica server IP address> 25000
```

If telnet client can connect to replica's port 25000 verify if there is a packet loss on the network link between the Master and Replica servers. High packet loss is also the main reason of file send and lost ACK errors, like

*EX00810 file sending aborted"*

*EX00830 lost acknowledgment"*

To check the packet loss rate run following test from the Master server:

- Open command line and run the following command

```
ping -l 64000 -n 100 <Replica IP address>
```

Check for any packet loss: Packets: Sent = 100, Received = 0, **Lost = 100 (100% loss)**

If there is a packet loss, then ping with 32000 packet size. ie

```
ping -l 32000 -n 100 <Replica IP address>
```

If still there is 100% loss, then continue to run ping test with packet size 16000, 8000 and 4000.

If you get 0% loss i.e Packets: Sent = 100, Received = 100, Lost = 0 (0% loss) on 32000 or 16000 or 8000 or 4000, then make below changes.

- Stop the scenario from RHA scenario management.
- Stop CA ARCserve RHA Engine service on both the Master and Replica servers
- Go to RHA installation folder : C:\Program Files\CA\ARCserve RHA\Engine. Open ws\_rep.cfg in a Notepad. Search for # ChunkLength = 64K. Change accordingly to the packet size which has 0% loss. For example, if there is 0% loss on packet size 16000 then change accordingly in ws\_rep.cfg

ChunkLength = 16k

**NOTE: Remove # in front of ChunkLength**

Save ws\_rep.cfg after making change on both the Master and Replica servers.

- Start CA ARCserve RHA Engine services.
- Start scenario from scenario management

**Note: Even a low % rate of packet loss can have a very significantly impact on replication performance and can bring replication to a complete stall.**

RHA uses a proprietary TCP/IP-based network protocol XONET. By default, it uses the following network settings:

TCPSendRecvBufferSize = 256K (TCP Window Size used by XONET)

Chunklength = 64K (buffer size used by XONET)

If the network latency between the Master and Replica servers exceeds 50ms RHA uses a lower Chunklength = 4K unless a different value is configured in ws\_rep.cfg file. KB articles below explain how to find the optimal Chunklength and TCPSendRecvBufferSize values:

<http://arcserve-knowledgebase.com/index.php?View=entry&EntryID=3493>

<http://arcserve-knowledgebase.com/index.php?View=entry&EntryID=599>

## Chapter 7: Slow Replication performance

During online replication, if the network throughput between the Master and Replica is not sufficient to keep up with the change rate on the Master, then the spool on the Master will be building up. When changes are queued in the spool on the Master and are not transferred to the Replica in real time it introduces a replication delay. Moreover, replication will completely stop if the spool disk runs out of space or the spool usage exceeds the configured threshold. You can see the Master and Replica spool usage in RHA GUI or RHA performance counters that show all important replication metrics such as:

- Update rate on the Master in replicated folders
- Network bandwidth used by replication
- Spool usage on the Master and Replica servers

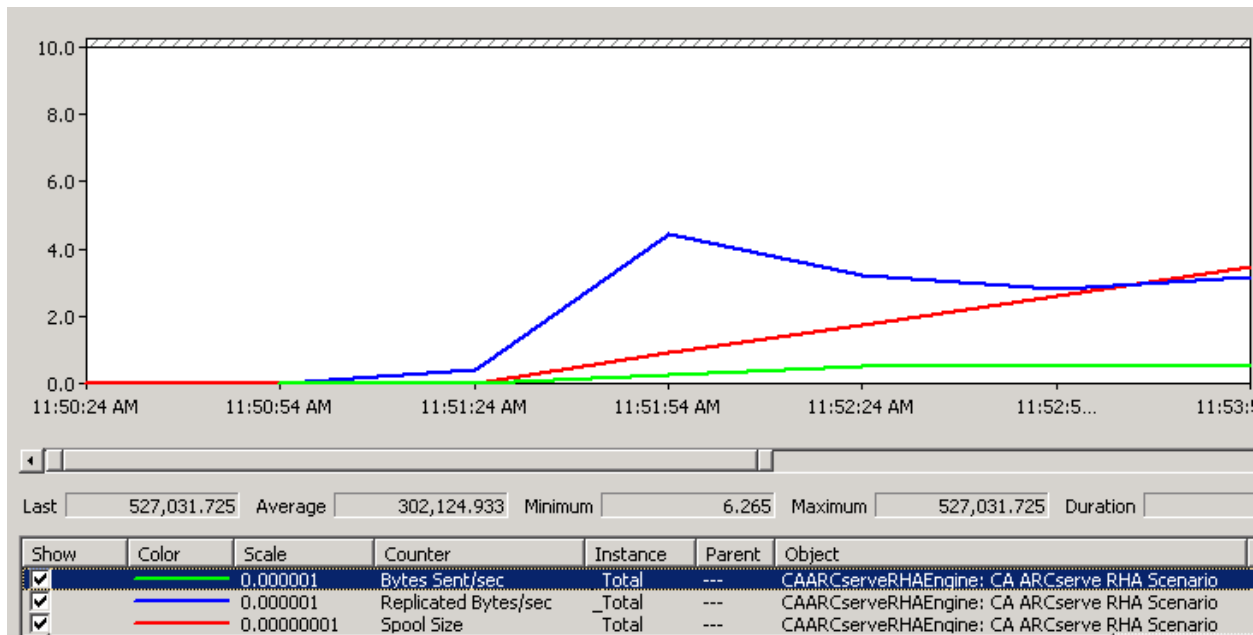
Open Windows performance monitor add the following counters:

Object: CAARCserveRHAEngine : CA ARCserve RHA Scenarios

Counters:

Bytes Send/sec (bandwidth used by replication)  
 Replicated Bytes/sec (update rate on the server)  
 Spool Size (space on disk used by spool on the server)

In the example below you see that Replicated Bytes/sec is much higher than Bytes Send/sec. As a result the scenario spool is continuously growing on the Master



**Note: If you don't see RHA performance objects after installing RHA Engine on a server follow the KB article below explaining how to restore RHA performance objects**

<http://arcserve-knowledgebase.com/index.php?View=entry&EntryID=3398>

The main reasons of spool growth on a Master server

- The Master server has extremely high I/O change rate (the update rate is much higher than the network bandwidth). Replication journals start queuing in spool on the Master causing significant spool growth and replication delay. When scenario is running highlight the Master server in "Scenario" window and click on "Statistics" tab in RHA GUI. In "Statistics" window RHA displays "Online file changes per root directory" that includes cumulative changes per root directory and in total since the scenario has been started (see the screenshot below). At the top of the window you can also see when the scenario was started. Based on this statistics data you can estimate the update rate. Divide total changes by number of seconds the scenario has been running and you'll obtain the average update rate.

Statistics

Master (Active) 'calabsrv05-ex2' Statistics

State	Running
Start of replication	04/17/13 19:06:28
Version	16.5.0.3409

Spool space:

Size	% of threshold
0 Bytes	0 %

Online file changes per root directory:

Root Directory	Size	Folders Created	Changed	Removed	Renamed
C:/Exchange/SG1/data	72 KB	0	2	0	0
C:/Exchange/SG1/log	1.2 MB	0	4	3	2
C:/Exchange/SG2/data	0 Bytes	0	0	0	0
C:/Exchange/SG2/log	22.5 KB	0	1	0	0
C:/Exchange/SG3/data	0 Bytes	0	0	0	0
C:/Exchange/SG3/log	22.5 KB	0	1	0	0
Total	1.32 MB	0	8	3	2

Transferred bytes to Replicas:

Host	Total Sent Data	Current File Name	Data To be Sent	Transmission Speed	Current Progress
calabsrv06-ex2	7.41 MB	0	0 Bytes	32 Bps	

To see how the update rate and the bandwidth consumed by replication changes in time you can create a performance log on the Master and Replica servers and start logging RHA performance counters. KB articles below explain how to create a performance log with RHA counters from Windows perfmon GUI or from a command line:

<http://arcserve-knowledgebase.com/index.php?View=entry&EntryID=10>

<http://arcserve-knowledgebase.com/index.php?View=entry&EntryID=3809>

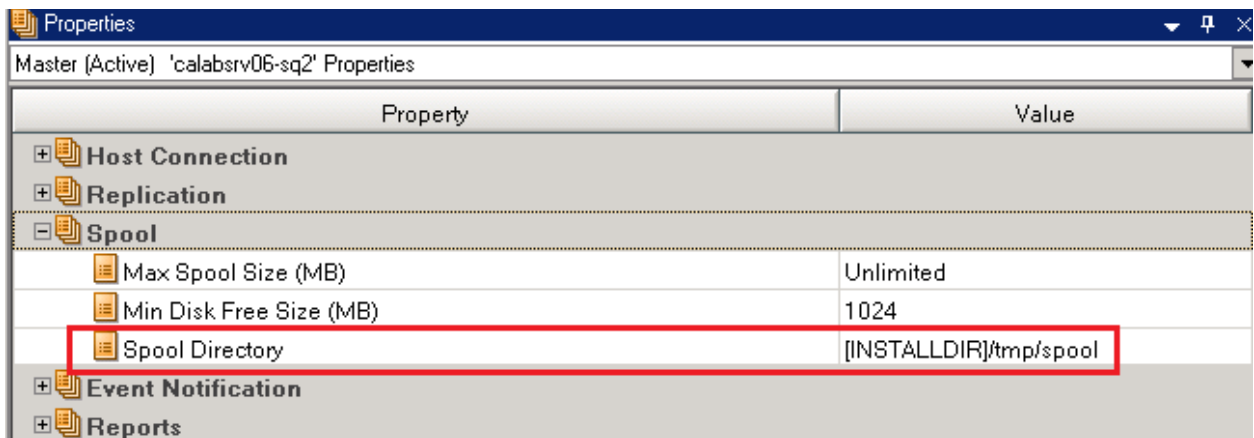
If the update rate is higher than the network throughput we recommend:

- Try reducing the amount of replicated data by excluding temporary files and other files/folders that don't need to be replicated
- When data is replicated over a slow WAN and the update rate is too high turn on **Compress Data During Transfer** in the Master server properties → Replication
- Increase **Number of Streams** in the scenario properties → Replication → Optional Settings. By default, the value is set to "User Configured" and RHA uses 1 stream if the network latency (RTT) between the Master and Replica does not exceed 50ms. If the network latency exceeds 50ms RHA uses 5 streams. You may benefit from using multiple streams even if the network latency is below

50ms. Setting the value explicitly forces RHA to use the configured number of streams.

**Note: It is not recommend using more than 1 stream on a LAN.**

- Check the network latency and bandwidth and try to find the most appropriate network settings following the KB articles below:  
<http://arcserve-knowledgebase.com/index.php?View=entry&EntryID=3493>  
<http://arcserve-knowledgebase.com/index.php?View=entry&EntryID=599>
- Upgrade your network to get more bandwidth and bring the network latency down.
- High rate of packet loss. Even 1-2% packet loss may cause a significant degradation of replication performance (see chapter 6).
- Spool is located on the same drive as the data being replicated causing a high amount of Disk I/O. The amount of I/O could also increases during full backups or maintenance jobs. Try configuring spool directory on a dedicated drive and exclude it from Antivirus and backup jobs. The default spool path is [INSTALLDIR]\tmp\spool. You can change the location of spool in the Master and Replica server properties:



**Note: Make sure that you have enough free space on the spool disk. The recommended free space is 15% of the replicated data size**

<http://arcserve-knowledgebase.com/index.php?View=entry&EntryID=4201>

- The spool disk has high I/O read and write latencies. When replication is running check the following performance counters:  
Physical Disk: Avg Disk Seconds/Read  
Physical Disk: Avg Disk Seconds/Write  
The averaged values for both counters should not exceed 5-10ms.
- Antivirus software scanning files written to the spool folder. Exclude RHA process and RHA spool folder from Antivirus real time protection and scans.  
<http://arcserve-knowledgebase.com/index.php?View=entry&EntryID=2494>

- The Master server is running low on system resources due to high CPU usage, shortage of memory, very high I/O rate or high kernel memory utilization. High kernel paged pool utilization is a very typical problem on a 32-bit OS particularly when 3GB flag is set in boot.ini. Running server low on system resources may significantly degrade replication performance and trigger I/O failures and replication errors. Follow recommendations in the KB articles below:

<http://arcserve-knowledgebase.com/index.php?View=entry&EntryID=2992>

<http://support.microsoft.com/kb/304101>.

## Chapter 8: Debugging Driver Errors

When starting a scenario you receive "Unable to Start File System Monitoring" error and in the Master log you see notifying errors message:

*Notifying scenario <scenario ID>, host 1; message id: CR00018, message: Unable to start file system monitoring*

The CA ARCserve RHA Engine uses a minifilter driver "XOMF" which is loaded on the Master when the scenario starts, and unloaded when the scenario stops.

If you receive the "Unable the start file system monitoring error", this usually indicates either a problem loading XOMF driver or attaching the driver to the replicated volume. To determine the root cause of this error, you can use the Microsoft Filter Manager utility (fltmc.exe). The driver is loaded only on a Master server. On the Master server that is having a problem do the following:

1. Open a command line window and run **fltmc** command. The output from this command shows a list of all loaded minifilter drivers. Check that **xomf** driver is not loaded.
2. Then try to load the driver manually:  
**fltmc load xomf**

Here are the most common errors and their resolutions:

***Load failed with error: 0x80070002***

***The system cannot find the file specified.***

While the file does exist in %WINDIR%\system32\driver folder on the Master, the most common reason is that the driver is not properly registered. See below the instructions how to re-register the driver:

**Windows 2003 (x86)**

To re-register the driver right click on INSTALLDIR\XOMF3\xomf3.inf file on the Master and select Install option.

**Windows 2003 (x64)**

To re-register the driver right click on INSTALLDIR\XOMF3x64\xomf3x64.inf file on the Master and select Install option.

**Windows 2008 (x86)**

To re-register the driver right click on INSTALLDIR\XOMF8\xomf8.inf file on the Master and select Install option.

**Windows 2008 (x64)**

To re-register the driver right click on INSTALLDIR\XOMF8x64\xomf8x64.inf file on the Master and select Install option.

Where INSTALLDIR is the directory where CA ARCServe RHA Engine is installed on Master that is by default:  
C:\Program Files\CA\ARCServe RHA\Engine

***Load failed with error: 0x8007052******A required privilege is not held by the client.***

1. Go to Administrative Tools -> User Rights Assignment -> Local Security Policy.
2. Under Load and Unload Drivers make sure the local administrator and the domain administrators group is listed.
3. Also check this Security Policy at the Domain level as Domain level policies supersede Local Security Policies.

***Load failed with error: 0x80070420******An instance of the service is already running.***

While "fltmc" does not show XOMF loaded, this indicates that the driver was not properly unloaded (partially unloaded). This may happen after upgrade of RHA Engine or after abnormal scenario stop. The only solution is to reboot the server.

After reboot you may need to re-register the driver using the driver INF file and try to load it manually.

If the command to load to driver "**fltmc load xomf**" does not return any error run "**fltmc**" command to check that XOMF driver in the list of the loaded filter drivers. For example, here is an output from "**fltmc**" command:

Filter Name	Num Instances	Altitude	Frame
xomf	2	303800	0
luafo	1	135000	0
RsFx0103	0	41001.03	0

Attach the driver to all volumes replicated in the scenario you are troubleshooting. For example, to attach XOMF driver to volumes C and D run the following commands

**fltmc attach xomf C:**

**fltmc attach xomf D:**

If the attach commands completed without errors you should be able to see all instances of XOMF by running the "fltmc instances" command:

```
xomf          C:          303800  xomf8x64 - Top Instance  0
xomf          E:          303800  xomf8x64 - Top Instance  0
```

If the driver loads successfully and can be attached to all replicated volumes, unload the driver by running "fltmc unload xomf" and try running the scenario again.

**Note: Make sure that the disk you are trying to attach the driver is online. This is particularly important for cluster and iSCSI drives that can be offline.**

When troubleshooting "Unable to start file system monitoring" errors check the preceding error messages in the Master log. For example, you may see the following errors:

***CR00077 Unable to mount intersecting root directories***

***CR00018 Unable to start file system monitoring***

You are most likely trying to replicate a volume that has a mount point for another volume you are also trying to replicate. You can check this by opening a command prompt on the Master and typing "mountvol". If you are trying to replicate the roots of F and G and the F drive has a mount point for G then you will get these errors because you are essentially trying to replicate part or all of G twice. Usually you can resolve these errors by removing one of the volumes from the scenario. In the above example, you can remove volume G from the scenario and allowing G to be replicated by just replicating F.

***CR00077 Unable to start file system monitoring - (invalid access to memory location)" or "FilterSendMessage failed, result: Invalid access to memory location"***

You start a scenario and see one of the errors above referencing we have an invalid access to memory location.

Apply Microsoft update 975759 and then reboot the server that received this error (generally the active server)

<http://support.microsoft.com/kb/975759>

***"CD00379" "Write to journal file failed <OS error code>"***

This is another critical error event stopping the scenario. The driver failed to write a journal file. Data integrity cannot be preserved; therefore the scenario must be stopped. The error code returned by OS is in hexadecimal format. In the Master log look for *FilterSendMessage* error sent in conjunction with *"Write to journal file failed"*. It usually shows the root cause of the failure. For example:

*FilterSendMessage XOMF\_RT\_JRN\_OP failed, result: **A device attached to the system is not functioning.***

or

*FilterSendMessage XOMF\_RT\_JRN\_OP failed, result: **The request could not be performed because of an I/O device error.***